Synchronous Hybrid Message-Adversary

Danny Dolev (HUJI)* and Eli Gafni (UCLA)†

regular submission, not eligible for the best student paper award

arXiv:1605.02279v2 [cs.DC] 5 Jul 2016

Abstract

The theory of distributed computing, lagging in its development behind practice, has been biased in its modelling by employing mechanisms within the model mimicking reality. Reality means, processors can fail. But theory is about predicting consequences of reality, hence if we capture reality by "artificial models," but those nevertheless make analysis simpler, we should pursue the artificial models.

Recently the idea was advocated to analyze distributed systems and view processors as infallible. It is the message delivery substrate that causes problems. This view not only can effectively emulate reality, but above all seems to allow to view any past models as *synchronous* models. Synchronous models are easier to analyze than asynchronous ones. Furthermore, it gives rise to models we haven't contemplated in the past. One such model, presented here, is the Hybrid Message-Adversary. We motivate this model through the need to analyze Byzantine faults. The Hybrid model exhibits a phenomenon not seen in the past.

*Email: danny.dolev@mail.huji.ac.il.

†Email: gafnieli@gmail.com.

1 Introduction

Recently, in [2], it was shown that a read-write wait-free asynchronous system can be modelled by a synchronous message passing system with a message adversary. This was a step beyond [5, 6] that equated shared-memory and message-passing when at most a minority of the processors are faulty. The work in [2], not only equated message-passing and shared-memory for the potential of $n-1$ faults, but in addition the message-passing system was synchronous. And, it paid handsomely with an almost trivial derivation of the necessity condition of asynchronous computability [23].

Here, as a first step, we extend the study of [2] to the m-resilient read-write asynchronous system, and find the message adversary m-AD in the synchronous message passing model that makes the synchronous system equivalent to the asynchronous one. m-AD is a system in which in each synchronous round m-AD chooses m processors and can remove any of the messages they send, subject to the condition that for the two messages exchanged between a pair of processors from the chosen m processors in the round, it removes at most one message.

This follows from the simple realization that in [2], when we take a system of two processors, instead of viewing the adversary as one that at each round can remove one message of the two sent, we could equivalently say that the adversary chooses a processor and can remove the message it sends. A simple semantic difference, that was not realized by us for a while.

The structure of 1-AD was realized 3 decades ago in [28] but not pursued beyond $m = 1$. The celebrated transformation of shared-memory to message-passing in [5, 6], which holds only for $m < (1/2)n$, is hard to cast as a message-adversary that chooses a set of processors and removes some of its messages. It is more than a decade later that the second author quoted in [4] discovered that after 3 rounds of the transformed system it can be viewed as a message-adversary.

When solving colorless tasks ([9]), if $m < (1/2)n$, we can do away with this last condition that prevents m-AD from removing the two messages exchanged by faulty processors. For the rest of the paper we consider only colorless tasks, hence we allow m-AD a complete freedom in removing in a round any messages from the m processors it chose for the round. We call this type of adversary *m-mobile omission faults*.

This paper deals with just a single colorless task. The colorless task we consider is the binary-vector-consensus-task. To explain this task we first put it within the historical perspective of the development of the various versions of the set-consensus task.

The ingenious original version of set-consensus task invented by Chaudhuri [13], was a version we call here, *set-election*: n processors start each with their own identifier (id) as its input, and each outputs a participating id, such that the number of distinct outputs should be smaller than n.

Using topological arguments this problem was later proved to be unsolvable read-write wait-free [7, 23, 27]. About a decade and a half later a technically simple but profound paper [3] formulated the task of *vector-set-consensus*. In the vector-set-consensus task, the n processors are faced with $n-1$ independent consensuses. Each processor outputs for at least one of the consensuses, such that, two processors that output for the same consensus satisfy validity and agreement for that consensus. A simple reduction to set-election showed the vector-set-consensus to be unsolvable read-write wait-free, and to be equivalent to set election.

But maybe, had we restricted each of the consensuses in the vector-set-consensus task to be just binary-consensuses, rather than multi-valued consensus, then the resulting binary-vector-set-consensus is solvable read-write wait-free? Precluding this possibility is equivalent to proving Generalized Universality [18], a major result, which succeeded only few years after the formulation of the vector-set-consensus task.

The various tasks above have the obvious analogue when we change $n-1$ to $k < n-1$. Indeed, Chaudhuri [13] proposed the problem in the context of m-resilient system, and asked whether the system can solve m-set consensus. The BG-simulation in [7, 9] showed that had her problem been solvable read-write m-resiliently, then set-election could have been solved for $m+1$ processors read-write wait-free, which is impossible. On the flip-side, a m-resilient system can trivially solve $m+1$ election!

But all these analyses were conducted in the context of benign failures. What if we have an asynchronous system with n processors and $f < (1/3)n$ Byzantine failures? Obviously, Byzantine failures are more serious failures than omission failures, and since we cannot solve f-binary-vector-set-consensus in the n processors f-resilient model we cannot solve it in the asynchronous n processors system with f Byzantine faults. Although the Byzantine failures are on a fixed set of processors, if these processors just lie about say their inputs, they cannot be detected, and a processor has to move on after receiving $n-f$ messages, lest the f it missed are the Byzantine set, which omitted messages. Hence $f+1$-set consensus is a lower bound in this case too. But can we solve the $f+1$-vector-binary-set-consensus problem, in n processors system with $f < (1/3)n$ Byzantine faults like we could do it for the analogue f-resilient system?

To our knowledge it is the first time this question has been asked. This is not surprising since the notion of binary-vector-set-consensus can only really be understood on the background of [18]. The question was asked with respect to the election version [14], but not for the vector-set-consensus task, let alone the binary-vector-set consensus task, since these tasks had not been introduced as yet, then. Why wasn't it asked with respect to the vector-set-consensus? Since these are reducible to each other without introducing the ideas in [18]. It looked meaningless to recast the results in [14] for vector-set-consensus. We do not believe that the general multi-valued set consensus task with its 24 (at the time!) possible versions, or for that matter even the general vector-set-consensus has any bearing on the vector-binary-set-consensus, when we consider Byzantine faults. Hence we do not pursue this multi-valued route, since we do not know a reasonable formulation of it in the Byzantine case. Thus, we are left with the question of the vector-binary-set-consensus power of asynchronous Byzantine n, f system.

We thus investigate this new general type of system. A *synchronous* system with f-fixed Byzantine faults and in addition m-mobile omission faults. We show this system to be equivalent in its vector-binary-set-consensus power to the same system of less severe failure of omission rather than Byzantine. We therefore investigate the general model of f-fixed omission faults and m-mobile omission faults. We call such a system that intermingles fixed omission failure and mobile omission failure, an *hybrid adversary*.

We completely characterize the vector-set-consensus power of the synchronous hybrid fix/mobile omission system with respect to all set-consensus versions, since in the non-hybrid benign faults all these versions are equivalent.

We do not show a reduction between the two synchronous systems one with the f-fixed Byzantine failures and m-mobile omissions, and the other with both types of failures being omission failures. Byzantine processors can always alter their input and behave correctly there on. Thus, when inputs conflict there is no resolution as to who are the "good ones" and who are the "bad ones." Rather, we should consider only contexts in which such resolution is not needed. To show such a result, the way we speculate to do it is first to show that the two systems have the same vector-binary-set-consensus power, which we do here.

Obviously the set consensus power of the benign system with fixed omissions and mobile omissions is stronger than the analogue synchronous hybrid Byzantine system. Thus a lower bound to the former is a lower bound to the latter. To show equivalence we present an algorithm for the Byzantine system that gets the same consensus power as that of the benign system, explicitly.

Our way of proceeding first with the binary-vector-set-consensus-power is in line with our recent think-

ing that systems that agree on set consensus tasks have the same power when other tasks are concerned. That is the thinking of: "set consensus tasks are the coordinates of any (reasonable) system," [15].

We obtain the result that an n, f, m omission system requires $n \geq (m+1)f + 1$ in order to solve the best value of set-consensus it can solve, $m + 1$. For lesser values of n it can solve set consensus $m + j$, when $n \geq \lfloor f(m+j)/j \rfloor + 1$.

This is the first time that we see a model where its set-consensus power changes gradually with n. In the m-resilient shared-memory case, we can always solve $m + 1$-set consensus. In the message passing system when $m > n/2 - 1$ we cannot solve anything. Once we are below that threshold we can solve $m + 1$-set consensus. It is the combination of faults that give rise to this gradually changing power phenomenon.

1.1 Outline of the Paper

We first show and prove the synchronous analogue of the asynchronous m-resilient model, and show that some restrictions on the adversary can be removed in the case of colorless-tasks ([9]).

We then characterize the hybrid n, f, m omission adversary for different combinations of these values. We use the [7] simulation to show that if a set of less than $(m + 1)f + 1$ processors can do $m + 1$ set consensus then $m + 2$ processors can do set-consensus wait-free contradicting [7, 23, 27].

Finally we show that the upper bound algorithm holds for the hybrid synchronous Byzantine system when dealing with binary-vector-set-consensus. For the upper bound we use the rotating coordinator algorithm [26, 12].

2 Problem Statement and Models

We assume a set of n processors $\Pi = \{p_1, p_2, ..., p_n\}$. The paper focuses on solving the k-vector-set consensus problem in the hybrid omission model SMPfm, to be explained below.

Definition 1 (k-vector-set consensus problem)**.** *Every processor has a vector of k initial inputs. Each processor, p, returns a vector $o_p[1..k]$ that satisfies:*

vset1: *If for any index $j \in [1..k]$ the inputs of all processors are the same, then every processor returns that value for index j.*

vset2: *For every entry $j \in [1..k]$ and for every two processors, p, and q, if $o_p[j] \neq \perp$ and $o_q[j] \neq \perp$ then $o_p[j] = o_q[j]$, and it is one of the inputs at entry j for some participating processor.*

vset3: *There exists an index $j \in [1..k]$ such that for every two processors, p, and q, $o_p[j] = o_q[j]$.*

Notice that requiring all to output for the *same* index j, is equivalent to asking just to output for some j, since processors can write their outputs, and then read and adopt outputs it sees. The first output written will be adopted by all.

Definition 2 (binary-k-vector-set consensus problem)**.** *Same as the above only that the vector of k initial inputs each processor has is a binary vector of 0's and 1's.*

The model SMPfm, called *hybrid omission* n, f, m, is a synchronous point-to-point message passing system. We consider an adversary that can remove messages. Before the start of the algorithm the adversary chooses a set S^f of f processors. In each round the adversary can remove some or all of the messages sent

by S^f. In addition in each synchronous round it can choose a set S^m of m processors and remove some or all of the messages sent from S^m. We call S^f, the fixed set, and S^m the mobile set.

Presenting models as "message adversary" enables us to easily deal with dynamic systems in which processors that presented an external erroneous behavior at one point to start behaving correctly later on, without the need to discuss what is their internal state when this happens.

Where is SMPfm coming from. The method in [2] transformed asynchronous wait-free SWMR shared memory into a synchronous message-passing adversary. This message passing adversary can be viewed, in hindsight, as an adversary that at each round chooses $n-1$ processors and removes some of their messages so that between two processors it chose, at least one of the two messages sent between them is not removed.

In the asynchronous m-resilient SWMR iterated shared memory, which is equivalent to the classical m-resilient non-iterated model [8, 21, 22], at each iteration processors are presented with a fresh SWMR memory initialized to $\perp$, and all write their cell and then snapshot the memory. The m-resiliency assumption entails that for each processor the snapshot returns at most m cells with the value $\perp$.

We now claim this model is equivalent to the synchronous message passing system in which a message adversary chooses a set S^m of m processors at a round and removes some of the messages they send. But, nevertheless, this removal is constrained by the condition that the adversary removes at most one of the messages sent between two members of S^m. We call this system constrained-SMPm.

Obviously an iteration of the asynchronous SWMR iterated shared memory, simulates a synchronous round of constrained-SMPm.

In the reverse direction we notice that constrained-SMPm can simulate an iterated SWMR collect step in $2n-1$ rounds [2], since a round of constrained-SMPm is obviously a round of constrained-SMP$_{n-1}$. Hence we can now simulate the Atomic-Snapshot algorithm in [1]. Since in each round of constrained-SMPm processors "read" from at least $n-m$ processors, the minimum size snapshot will be $n-m$.

Can we remove the constrain on constrained-SMPm, so that the message adversary chooses S^m and can possibly remove all their messages violating the requirement that among pairs in S^m at least one message survives? We can, when all we want is to solve colorless tasks and $m < (1/2)n$. When $m < (1/2)n$ a processor that through the rounds hears that at least $n/2+1$ processors have heard from it, can be sure that in the next round all will, since one of these $n/2+1$ processors will not be chosen by the adversary in the next round. Hence we can see that at least $n-m$ processors can progress simulating read write. What we lose is that if the adversary sticks with a fixed set S^m those processors cannot communicate. But at least $n-m$ processors will have outputs and S^m can adopt theirs.

Our last system of interest is SBMPfm. This system is like SMPfm only that the adversary can now tamper with messages from the fixed set S^f. Thus, it is, in a sense, like having f fixed byzantine faults and m mobile omission faults.

3 The Lower Bound

Theorem 1. *For the hybrid omission n, f, m, if $n < f(m+1)+1$, there is no algorithm to solve $(m+1)$-vector-set-consensus.*

Proof. The proof is based on a simulation that uses constructions similar to those used in [19, 20]. The details appear in the references, but nevertheless we sketch the construction of the simulation.

W.l.o.g. by way of contradiction an algorithm exists with all processors sending messages to all. Take

$m+2$ BG processors [19, 20] simulating one round before moving to the next. In each round a simulator decides by safe-agreement [7, 9] whether a message sent is received or not. A BG processor will claim that a message was removed if it does not know the simulated local state of the processor that sends the message.

Since we are dealing with SMPfm we can equivalently deal with the election-version of set-consensus. Initially, every simulator tries to install its input value as the input to all simulated processors. At most $m+1$ safe agreements modules may be blocked and the corresponding processors cannot be simulated as sending messages. Thus, to proceed, when a simulator does not know the local state of a processor, then it will try to reach agreement that the message was removed.

But in the meantime, the safe agreement for this processor might be resolved, hence other processors may contend that a message was sent. If we do these message delivery safe agreements for a processor proceeding negatively (a message was not sent) from the highest index receiver to the lowest, and in the opposite positive direction, from low index to high, for the message sent, at most one safe-agreement about a single message from this processor may be blocked at the index that the positives and the negatives meet. This will manifest itself as a message to some processor that we do not know whether it was removed or received, and therefore we do not know the local state of this potentially receiving processor.

Thus, the initial possible lack of knowledge about the $m+1$ inputs may propagate to at most $m+1$ omission failures from round to round. Thus we get an execution of a synchronous system $n, m+1$ with n processors and $m+1$-omission faults.

Now we move to the ramification of this simulation. If there exists an algorithm for the system n, f, m such that each run of n, m can be viewed as a run of n, f, m, then the algorithm that obtains $m+1$-vector-set consensus for n, f, m will be an algorithm for $m+1$-vector-set consensus for the system $n, m+1$, contradicting [7, 23, 27].

Hence, it must be that there exits a run of $n, m+1$ that cannot be explained as a run of an algorithm for $m+1$-vector-set consensus of a synchronous hybrid system n, f, m.

The system n, f, m can have $m+1$ omission failure in a round by choosing each round $m+1$ processors from which to remove messages as to simulate a round of $n, m+1$. Observe that at least one of these $m+1$ processor chosen at a round has to be on the account of the f fixed processors, since in n, f, m we have only m-mobile omission faults at a round rather than $m+1$.

Thus, we want to show that for n small enough, $n < f(m+1)+1$, given a run of $n, m+1$ we can allocate f-fixed processors that explain the run as a run of n, f, m.

We take an infinite run of $n, m+1$ and divide it into chunks of n rounds. In the first round processors $1, 2, \ldots, m+1$ omit, in the second round $2, 3, \ldots m+2$ omit, etc, with wrap around at round n. Then repeat the same for the second chunk etc.

Thus, if we take any k chunks like this, a processor appears in all the chunks $(m+1)k$ times each at a different round. Thus, since we have f fixed faults, the largest number of rounds we can justify with these fixed faults is $f(m+1)k$. But if $nk > f(m+1)k$ we will not be able to justify it as an n, f, m run.

Obviously for any $n \leq f(m+1)$ we will be able to justify it as an n, f, m run. Our example made the repetition of each processor equal. If it is unequal we will attribute the processors that are at the top f in ranking of repetitions as the fixed set and will able to justify any run of $n, m+1$ as a run of n, f, m. □

The lower bound proof implies that if we take $n \geq f(m+1)+1$ we will not be able to explain the run as an $n, m+1$ run. What is left is to show that indeed for $n \geq f(m+1)+1$ we have an algorithm that obtains $(m+1)$-vector-set consensus.

An easy repetition of the lower bound arguments above show that if $n < \lfloor f(m+j)/j \rfloor + 1$ we cannot

solve $m + j$-set consensus.

4 Binary-k-vector-set Consensus

For simplicity of exposition we will use only the binary-k-vector-set consensus even for SMPfm as we know we can solve the multi-valued one using [18]. To show the phenomenon of consensus power growing gradually with n given fixed f and m, we assume for convenience that $m + f < n/2$ and $f < n/4$. The algorithm rely on the idea of the rotating coordinator [26, 12]. Only that in hindsight we know that each phase of the implementation hides a COMMIT_ADOPT protocol [17]. [1] Hence, we first pause the presentation and show how we solve COMMIT_ADOPT in SMPfm and then in SBMPfm.

4.1 Commit-Adopt implementation for SMPfm, and SBMPfm

In the COMMIT_ADOPT protocol each processor invokes COMMIT_ADOPT with an initial value. Each processor, p, returns as an output a pair $\langle v_p, e_p \rangle$, where $e_p \in \{\text{COMMIT}, \text{ADOPT}\}$. COMMIT_ADOPT ensures that:

CA1: If all processors invoke COMMIT_ADOPT with the same value, then every processor, p, returns with that value and with $e_p = \text{COMMIT}$.

CA2: If a processor, p, returns with $e_p = \text{COMMIT}$, then for any processor q, $v_q = v_p$.

A COMMIT_ADOPT algorithm in [17] is given in a wait-free shared-memory model. To use this algorithm in different models we either implement the shared-memory in the model, or just show an implementation that comply with the properties that make COMMIT_ADOPT algorithm in shared-memory work. The properties are to have two iteration where at most one value will be observed as a proposal to commit in the second round, and if a processor views only commit proposal in the second round, then any other processor in the second round will observe a proposal to commit.

In SMPfm if all processors start with the same bit, every processor will get all messages of the same bit and there will be at least $n/2 + 1$ of them. A processor that receives all same bit, will propose in the second round to commit that bit. Obviously, since majority sent same bit, no other processor will propose to commit a different bit.

In the second round, a processor commits if it obtains at least $n/2+1$ proposals to commit. A processor that does not propose to commit does not send a message. Obviously if one processor receives at least $n/2 + 1$ proposals to commit, others cannot miss all these proposals and will see at least one proposal to commit, and hence will adopt that bit.

In SBMPfm we have to worry about the f processors whose messages the adversary is allowed to tamper with. Thus, we cannot require to see all messages of the same bit, since then the adversary can prevent commit in the case that all started with the same bit. Nevertheless we know that if all started with the same bit then a processor will get at least $n/2 + 1$ of that bit, and at most f of the complement bit.

Hence, we use this test in order to propose commit at the second round. Since in the worst case the complement bit was send by "correct" processors, we nevertheless are left with more than quarter of correct processors whose input is that bit. Hence, this set of processors in the first round will prevent any other processor to propose to commit the complement bit, as the size of the set is greater than f. A processor that does not propose to commit does not send anything in the second round.

[1] The COMMIT_ADOPT protocol is the concept behind last two rounds of the original Gradecast [16] protocol.

In the second round for a processor to "read" a commit proposal it is enough if it obtains at least $n/4$ commit proposals of the same bit. To commit, a processor needs to receive again at least $n/2+1$ proposals of commit of the same bit (We can now ignore commit of the different bit since there can be anyway at most f of them). Again we can argue that in the second round, if any processors commits, we must have at least $n/4$ correct processors which send commit proposal of that bit and hence all will at least adopt that bit.

Figure 1 and Figure 2 present both versions of the COMMIT_ADOPT protocol.

Algorithm 1: COMMIT_ADOPT(v_p): The Commit Adopt protocol for omission faults

1: **round 1:** send v_p to all; */* executed by processor p */*

2: **round 2: if** all "bits" received are the same **then send** the "bit" to all **else do not send**;

3: */* Notations */*

4: **let** maj be the bit received and let $\#maj$ be the number of processors that sent it;

5: */* Scoring */*

6: **if** $\#maj > n/2$ **then set** e_p :=COMMIT **else set** e_p :=ADOPT;

7: **if** $\#maj > 0$ **then set** $v_p := maj$; */* otherwise remain with the original v_p */*

8: **return** $\langle v_p, e_p \rangle$.

Algorithm 2: COMMIT_ADOPT(v_p): The Commit Adopt protocol for Byzantine faults

1: **round 1:** send v_p to all; */* executed by processor p */*

2: */* Notations */*

3: **let** maj be the bit received the most and let$\#maj$ be the number of processors that sent it;

4: **let** min be the bit received the least and let its number be $\#min$;

5: **round 2: if** $\#min \leq f$ and $\#maj > n/2$ **then send** maj **else do not send**;

6: */* Notations */*

7: **let** maj be the bit received the most and let $\#maj$ be the number of processors that sent it;

8: */* Scoring */*

9: **if** $\#maj > n/2$ **then set** e_p :=COMMIT **else set** e_p :=ADOPT;

10: **if** $\#maj \geq n/4$ **then set** $v_p := maj$; */* otherwise remain with the original v_p */*

11: **return** $\langle v_p, e_p \rangle$.

4.2 Binary-k-vector-set Consensus Protocol

We first focus on the case of $k = m+1$. For the discussion below assume that $n = f(m+1)+1$. The idea of the protocol is to run in parallel the basic process for each of the $m+1$ entries in the binary-$(m+1)$-vector-set consensus. The process below will ensure that in each entry in the output vector different processors never produce conflicting outputs, and that for at least one entry all processors report an output.

We assign $m+1$ coordinators to each phase of the protocol, one per entry in the vector. The coordinators play a role in a specific round of sending messages in each phase, as described below. We run the protocol for $f(m+1)+2$ phases, each takes three rounds of message exchange.

For a given entry all processors repeatedly exchange their values in each phase. Each phase begins with concurrently running a COMMIT_ADOPT on the current values of all processors. In the first phase processors

use their initial input values, and later phases the values computed by the end of the previous phase.

Following the COMMIT_ADOPT step the coordinator of the current phase broadcasts the value it obtained from the recent COMMIT_ADOPT.

A processor that completed the recent COMMIT_ADOPT with COMMIT ignores the coordinator's message and updates its value to be the committed value of the COMMIT_ADOPT. A processor that did not complete the recent COMMIT_ADOPT with COMMIT adopts the value it receives from the coordinator, if it received a value, if no value was received it remains with its original value.

By the end of this value updating we are guaranteed that if the coordinator was correct when it sent its coordinator's value, then all processors will end up holding identical values. The reason is that the COMMIT_ADOPT properties imply that if a processor returns from the COMMIT_ADOPT with COMMIT, all processors return from the COMMIT_ADOPT with identical values, so this is also the value the correct coordinator sends. If this is not the case, every processor adopts the coordinator's value, and again they hold identical values.

Observe that our assumptions are that all processors receive the values from all correct processors, even when the adversary chooses to change their messages. Therefore, the current coordinator received the correct value from the COMMIT_ADOPT as every other processor.

Once all processors hold identical values, in all future phases the COMMIT_ADOPT at each processor will return COMMIT with that value, no matter who the rest of the coordinators are.

The above basic process is repeated for $f(m+1)+2$ phases for all the $(m+1)$ entries of the $(m+1)$-vector-set consensus. After the end of the last phase each processor reports output for every entry in which the latest COMMIT_ADOPT returned COMMIT. The COMMIT_ADOPT properties imply that there will not be any conflict on output values in any index. Moreover, for each entry for which in one of the first $f(m+1)+1$ phases there happened to be a correct coordinator sending its value, all processors return that value for that entry.

What we are left to discuss is why there would always be at least one correct coordinator in at least one entry in at least one phase. Although this argument is repetition of the argument in the lower-bound section, we repeat it here. Observe that we assign to each phase $m+1$ different coordinators. The assignment of coordinators to phases is such that for $n = f(m+1)+1$ each one appears in exactly $f+1$ different phases. This implies that there can be at most $f(m+1)$ phases in which at least one of the coordinators assigned to entries in that phase is from the fixed set f. Look at a phase in which no coordinator is from the fixed set. The adversary can drop messages from at most m of the coordinators that send their coordinators' values in that phase. Therefore, there should be an entry at which the coordinator sending the coordinator's value is correct.

Observe that for binary values one can replace the condition in Line 10 of Figure 2 to $\#maj > 0$, since if no processor returns with COMMIT then non-Byzantine processors have sent both 0 and 1. For non-binary values, instead of testing for $\#min$ we need to test for non-max values, and can replace the condition in Line 10 of Figure 2 to $\#maj > f$. Moreover, one can add a filtering in Line 3 to filter out values that do not conform with what one expects to receive, since they are clearly being sent by Byzantine processors. Similar filter can be used in Line 9 of Figure 3. We do not have any use for such a filtering in the protocols of the current paper.

One can generalize the lower bound proof of Section 3 for $m+k$-vector-set consensus algorithm for $k \geq 1$ to obtain a lower bound of $n > \lfloor \frac{f(m+k)}{k} \rfloor$. The DYNAMIC $(m+1)$-VECTOR-SET CONSENSUS protocol of Figure 3 can be changed accordingly and will run in $\lfloor \frac{f(m+k)}{k} \rfloor + 2$ phases. When increasing the number of of processors by m one can device a protocol that runs for only $\lfloor \frac{f}{(k-m)} \rfloor + 2$ phase. Thus for

Algorithm 3: DYNAMIC $(m+1)$-VECTOR-SET CONSENSUS:
a *Byzantine* $m+1$-vector-set consensus algorithm

```
1:  /* Initialization */                                   /* executed by processor p*/
2:  let v(j) be the initial input to consensus index j, 1 ≤ j ≤ m+1;   /* the input values*/
3:      /* the permutation over the set of n processors */
4:  let s_{i,j} = p_ℓ, where ℓ = (i+j-1) mod n, for 1 ≤ i ≤ f(m+1)+2, 1 ≤ j ≤ m+1;

5:  /* Main loop for each of the k indices, 1 ≤ j ≤ m+1, all of them in parallel */
6:  for phase i := 1 to f(m+1)+2 do
7:      ⟨v̂(j), ê(j)⟩ =COMMIT_ADOPT(v(j));

8:      if p = s_{i,j} then send v̂(j) to all;   /* the rotating coordinator for index j sends its value */
9:          let v'(j) be the value received from s_{i,j};   /* ⊥ if no value was received */

10:     if ê(j) = COMMIT then v(j) := v̂(j)
11:         else if v'(j) ≠ ⊥ then set v(j) := v'(j);   /* adopt coordinator s_{i,j} value*/
12: end for

13: for each j, 1 ≤ j ≤ m+1: if ê(j) = COMMIT then o_p(j) := v(j) else o_p(j) := ⊥;

14: return o_p[1..m+1].
```

$n \geq f(m+1)+1$ we can solve $m+1$-set consensus. for $n \geq f(m+2)/2+1$ we can solve $m+2$ set consensus, etc.

All the formal proofs appear in the appendix.

5 Conclusions

We introduced a new type of distributed-system call Hybrid-Message-Adversary. It gives rise to phenomenon never seen before of set-consensus power changing gradually even though the various types and number of faults do not change. In our mind the only notion of set consensus that makes sense in the Byzantine setting is that of binary-vector-set-consensus. To our knowledge we are the first to ask this question, and in fact we are still at loss but not far, we suspect, from an answer.

Next, we can imagine message adversary with mobile Byzantine faults and combinations thereof with omission fixed or mobile faults etc.. In fact, the analogue of message adversary with mobile Byzantine faults was studied in the domain of Cryptography under the name of mobile viruses, transient or proactive faults [24, 10, 25, 11], but none looked at the relative power of tasks, let alone the set-consensus power.

Why should we? We recently [15] started to suspect that "natural systems" can be characterized by their set consensus power. Thus if this is proved and we equate the set-consensus power of synchronous Byzantine of f faults and SBMPfm with $m = f$, then they will be equivalent.

References

[1] Yehuda Afek, Hagit Attiya, Danny Dolev, Eli Gafni, Michael Merritt, and Nir Shavit. Atomic snapshots of shared memory. *J. ACM*, 40(4):873–890, September 1993.

[2] Yehuda Afek and Eli Gafni. A simple characterization of asynchronous computations. *Theor. Comput. Sci.*, 561:88–95, 2015.

[3] Yehuda Afek, Eli Gafni, Sergio Rajsbaum, Michel Raynal, and Corentin Travers. The k-simultaneous consensus problem. *Distributed Computing*, 22(3):185–195, 2010.

[4] H. Attiya and J. Welch. *Distributed Computing: Fundamentals, Simulations and Advanced Topics*. John Wiley & Sons, 2004.

[5] Hagit Attiya, Amotz Bar-Noy, and Danny Dolev. Sharing memory robustly in message-passing systems. In *Proceedings of the Ninth Annual ACM Symposium on Principles of Distributed Computing*, PODC '90, pages 363–375, 1990.

[6] Hagit Attiya, Amotz Bar-Noy, and Danny Dolev. Sharing memory robustly in message-passing systems. *J. ACM*, 42(1):124–142, January 1995.

[7] Elizabeth Borowsky and Eli Gafni. Generalized FLP impossibility result for t-resilient asynchronous computations. In S. Rao Kosaraju, David S. Johnson, and Alok Aggarwal, editors, *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, May 16-18, 1993, San Diego, CA, USA*, pages 91–100. ACM, 1993.

[8] Elizabeth Borowsky and Eli Gafni. A simple algorithmically reasoned characterization of wait-free computations (extended abstract). In *Proceedings of the Sixteenth Annual ACM Symposium on Principles of Distributed Computing, Santa Barbara, California, USA, August 21-24, 1997*, pages 189–198, 1997.

[9] Elizabeth Borowsky, Eli Gafni, Nancy A. Lynch, and Sergio Rajsbaum. The BG distributed simulation algorithm. *Distributed Computing*, 14(3):127–146, 2001.

[10] Ran Canetti and Amir Herzberg. Maintaining security in the presence of transient faults. In Yvo G. Desmedt, editor, *Advances in Cryptology — CRYPTO '94: 14th Annual International Cryptology Conference Santa Barbara, California, USA August 21–25, 1994 Proceedings*, pages 425–438, Berlin, Heidelberg, 1994. Springer Berlin Heidelberg.

[11] Miguel Castro and Barbara Liskov. Proactive recovery in a byzantine-fault-tolerant system. In *Proceedings of the 4th Conference on Symposium on Operating System Design & Implementation - Volume 4*, OSDI'00, pages 19–19, Berkeley, CA, USA, 2000. USENIX Association.

[12] Tushar Deepak Chandra and Sam Toueg. Time and message efficient reliable broadcasts. In *Proceedings of the 4th International Workshop on Distributed Algorithms*, pages 289–303, New York, NY, USA, 1991. Springer-Verlag New York, Inc.

[13] S. Chaudhuri. More choices allow more faults: Set consensus problems in totally asynchronous systems. *Information and Computation*, 105(1):132 – 158, 1993.

[14] R. De Prisco, D. Malkhi, and M. Reiter. On k-set consensus problems in asynchronous systems. *Parallel and Distributed Systems, IEEE Transactions on*, 12(1):7–21, Jan 2001.

[15] Carole Delporte-Gallet, Hugues Fauconnier, Eli Gafni, and Petr Kuznetsov. A complete characterization of set-consensus collections. Technical report, 2016. In sumbmission.

[16] Paul Feldman and Silvio Micali. Optimal algorithms for byzantine agreement. In *STOC '88: Proceedings of the twentieth annual ACM symposium on Theory of computing*, pages 148–161, 1988.

[17] Eli Gafni. Round-by-round fault detectors (extended abstract): Unifying synchrony and asynchrony. In *Proceedings of the Seventeenth Annual ACM Symposium on Principles of Distributed Computing*, PODC '98, pages 143–152, 1998.

[18] Eli Gafni and Rachid Guerraoui. Generalized universality. In Joost-Pieter Katoen and Barbara König, editors, *CONCUR 2011 - Concurrency Theory - 22nd International Conference, CONCUR 2011, Aachen, Germany, September 6-9, 2011. Proceedings*, volume 6901 of *Lecture Notes in Computer Science*, pages 17–27. Springer, 2011.

[19] Eli Gafni, Rachid Guerraoui, and Bastian Pochon. From a static impossibility to an adaptive lower bound: the complexity of early deciding set agreement. In Harold N. Gabow and Ronald Fagin, editors, *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 714–722. ACM, 2005.

[20] Eli Gafni, Rachid Guerraoui, and Bastian Pochon. The complexity of early deciding set agreement. *SIAM J. Comput.*, 40(1):63–78, 2011.

[21] Eli Gafni and Sergio Rajsbaum. Distributed programming with tasks. In *Principles of Distributed Systems - 14th International Conference, OPODIS 2010, Tozeur, Tunisia, December 14-17, 2010. Proceedings*, pages 205–218, 2010.

[22] Maurice Herlihy and Sergio Rajsbaum. The topology of shared-memory adversaries. In *Proceedings of the 29th Annual ACM Symposium on Principles of Distributed Computing, PODC 2010, Zurich, Switzerland, July 25-28, 2010*, pages 105–113, 2010.

[23] Maurice Herlihy and Nir Shavit. The topological structure of asynchronous computability. *J. ACM*, 46(6):858–923, November 1999.

[24] Rafail Ostrovsky and Moti Yung. How to withstand mobile virus attacks (extended abstract). In *Proceedings of the Tenth Annual ACM Symposium on Principles of Distributed Computing*, PODC '91, pages 51–59, 1991.

[25] Tal Rabin. A simplified approach to threshold and proactive RSA. In Hugo Krawczyk, editor, *Advances in Cryptology — CRYPTO '98: 18th Annual International Cryptology Conference Santa Barbara, California, USA August 23–27, 1998 Proceedings*, pages 89–104, Berlin, Heidelberg, 1998. Springer Berlin Heidelberg.

[26] Rdiger Reischuk. A new solution for the byzantine generals problem. *Information and Control*, 64(13):23 – 42, 1985. International Conference on Foundations of Computation Theory.

[27] Michael Saks and Fotios Zaharoglou. Wait-free k-set agreement is impossible: The topology of public knowledge. *SIAM Journal on Computing*, 29(5):1449–1483, 2000.

[28] Nicola Santoro and Peter Widmayer. Time is not a healer. In B. Monien and R. Cori, editors, *STACS 89: 6th Annual Symposium on Theoretical Aspects of Computer Science Paderborn, FRG, February 16–18, 1989 Proceedings*, pages 304–313, Berlin, Heidelberg, 1989. Springer Berlin Heidelberg.

Appendix

A Proofs

We prove the correctness of COMMIT_ADOPT for the SBMPfm model. The proof for the SMPfm is much simpler.

Claim 1. *Assume $n > \max(2f + 2m, 4f)$ and SBMPfm adversary. The protocol of Figure 2 meets the* COMMIT_ADOPT *requirements.*

Proof. To prove property CA1 observe that if all processors send an identical value in Line 1 of Figure 2 then each processor can receive at most f different values, from processors the adversary tampers with their messages. Therefore, at every processor the test in Line 5 will bring it to propose this identical value. Exactly for the same reasons each processor, p, will set up e_p to be COMMIT and will return the identical value.

For CA2, assume that processor p proposes a value in the second round of the protocol. Thus, the test in Line 5 is $true$. Denote by maj_p, $\#maj_p$ and $\hat{n}_p$ the parameters processor p used in Line 5 and V_p the multiset of $\bar{n}_p$ values it received, and respectively for another arbitrary processor q. Consider two cases. The first case is when $m \geq f$. In our model V_p may not contain at most the m values from omission faults and some of the values from the Byzantine processors. Since $\#maj_p > n/2$ and $\#maj_p \geq \hat{n}_p - f$ there are at least $m + 1$ values in V_p from processors that are correct at the current sending step and sent v_p to all. V_q should contain all these processors' values. The assumption that $m \geq t$ implies that V_q contains at least $f + 1$ value of v_p, therefore, Line 5 of Figure 2 implies that if decides to propose a value it should be that $v_q = v_p$. Now consider the case that $m < f$. In this case the condition $n > 4f$ implies that $n/2 > 2f$. This implies that V_p contains at least $n/2 - f$ values from processors that are correct at the current sending step that sent v_p, thus at least $f + 1$, and the rest of the above arguments hold.

The above argument implies that no two non-Byzantine processors send different values in the second round. Assume that processor p commits in Line 9. Let v be a value committed to. Assume the above notations for the messages received in the 2nd round. Neither V_q nor V_p can contain more than f non $\perp$ values the are not v. Thus, the protocol implies that $maj_p = v$. By definition we know that $|V_q \cap V_p| \geq n/2$, all of which are correct at the current message sending step. Since $\#maj_p \geq \hat{n}_p - f$ and $\#maj_p > n/2$, V_q contains at least $n/2 - f$ copies of v, thus more than $n/4$ and more than f. □

Claim 2. *Assume $n > \max(2f+2m, 4f)$ and SBMPfm adversary. In the* DYNAMIC $(m+1)$-VECTOR-SET CONSENSUS *protocol, if all processors' initial values to a given index j are the same, they output that value for that index at the end.*

Proof. In DYNAMIC $(m + 1)$-VECTOR-SET CONSENSUS processors update their initial values either in Line 10 or Line 11 of Figure 3. By Claim 1 we know that any invocation of it will result with all processors obtaining the identical value and with evaluation "commit", therefore Line 11 will never be executed, and at the end of the protocol all will produce an identical value. □

Claim 3. *Assume $n > \max(2f+2m, 4f)$ and SBMPfm adversary. In the* DYNAMIC $(m+1)$-VECTOR-SET CONSENSUS *protocol, if for some index j at some phase the coordinator p is correct when executing Line 8, then after executing Line 10 or Line 11 of Figure 3 in that phase, $v(j)$ is identical at all processors.*

Proof. Let p be the correct processor executing Line 8 of some index j, in some phase r. Consider two cases. If there is any processor that completes the COMMIT_ADOPT in Line 7 of phase r with "commit" and the other if none. In the first case, by the COMMIT_ADOPT properties we know that all processors complete the COMMIT_ADOPT with the same value (including the faulty processors). Therefore, the value p will send in Line 8 is the same, and therefore all processors will complete Line 9 with the same value. This implies that no matter which of the two lines, Line 10 or Line 11, any processor executes, all obtain the same value. In the second case, no processor completes Line 7 with "commit", and therefore all will execute Line 11 and will obtain the value the correct sender sent when it was correct while executing Line 8. □

Theorem 2. *If* $n > \max\big(f(m+1), 4f, 2(m+f)\big)$ *and assuming SBMP*fm *adversary, the* DYNAMIC $(m+1)$-VECTOR-SET CONSENSUS *protocol satisfies the properties of* k*-vector-set consensus, for* $k = m+1$.

Proof. In order to prevent having any correct coordinator in a phase all $m+1$ processors assigned to be coordinators in that phase need to either be in the fixed set of f faulty, or one of the m processors that suffer from omission in that phase. The definition of $s_{i,j} = p_\ell$, where $\ell = (i+j-1) \mod n$, for $1 \leq i \leq f(m+1)+2$, $1 \leq j \leq m+1$, assigns each processor p to exactly $m+1$ times in the first $f(m+1)+1$ phases. The fix set of f processors appear in at most $f(m+1)$ phases. Therefore, there is a phase in which none of them appear. Since omission faults can silence at most m processors in that phase, there is a correct processor executing Line 8 of Figure 3 in that phase. Let j be the index of that processor.

By Claim 3 we know that by the end of that phase all processors hold the same values in their $v(j)$. From the next phase on, until the end of phase $f(m+1)+2$ all will complete Line 7 in Figure 3 with "commit", and will end up having the same value in the j-th index. Moreover, for any other index j', two processors that assign a value to that index, assign the same value, since it is the value they completes Line 7 of index j' with "commit", and it is an identical value. The remaining property of DYNAMIC k-VECTOR-SET CONSENSUS obviously holds as well. □